Creative Collaboration

General Examination Administered by Mitch Resnick

Benjamin Mako Hill
mako@mit.edu

July 16, 2010

OBSTACLES TO LARGE GROUP CREATIVITY

> *What are the most significant obstacles and challenges that confront large numbers of people working together on creative design activities? How can technologies reduce or ease these challenges?*

Introduction

The rapid diffusion of systems such as email and discussions boards, instant messaging, wikis, and social media, has facilitated collaboration among previously unprecedentedly large groups. For example, English Wikipedia integrates the work of hundreds of thousands of contributors each month and today represents the aggregate product of millions of individuals' work (Ortega, 2009). Facebook now involves more than a billion users engaged in collaborative information gathering and distribution, sensemaking (Weick, 1995), and social interaction. Making sense of this phenomena is difficult enough. The challenges to supporting and reproducing the most successful of these examples are overwhelming.

In 1994, HCI and CSCW researcher Jonathan Grudin published an influential article on *Groupware and social dynamics: Eight challenges for developers* (1994b). In the article, Grudin lists a series of major obstacles to the development and diffusion of technological collaboration support tools as he understood them at that time. Grudin's article included

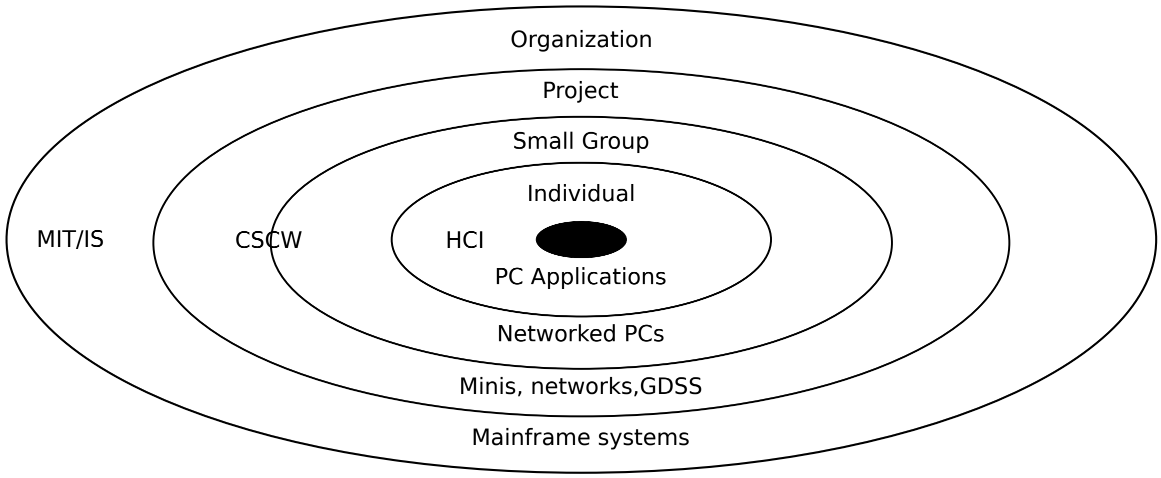

Figure 1: A partial reproduction of Figure 1 from Grudin (1994a,b) showing "development and research contexts" in the academic study of computer use in computer science. On the left side are the sub-fields or research streams in computer science. On the top are the types of user being served. On the bottom are the types of products being produced.

a figure – also published, slightly modified, in Grudin's (1994a) article on the history of CSCW – that describes the computing and research landscape along a spectrum from research focused on individuals (i.e., HCI), to groups (i.e., CSCW and Groupware), to organizations (i.e., MIS and IT). Grudin's diagram is reproduced in Figure 1.

Large-scale, supra-organizational collaborations, like the kind I have described, seem like they would naturally occupy a larger concentric ring in Grudin's diagram. These projects involve more people and larger, more complex, goals and outputs. If this were the case, one might then expect that MIS and IT researchers would have the most to say about these phenomena. But this has not been the case. CSCW researchers, historically focused on small group interaction and teams (as shown in the diagram), have led the academic charge into social computing and mass collaboration more broadly. The 2010 *Proceedings of CSCW* included a large number of papers on Wikipedia and others on social media.[1] Some of the most well respected CSCW and Groupware researchers have assumed new roles as respected researchers of social media and large-group collaboration.

For example, HCI pioneer Ben Shneiderman has published on social media and social

[1] http://www.cscw2010.org/program/papersnotes.php

computing more broadly (2007; 2008). CSCW pioneer Bonnie Nardi and her colleagues have published seminal studies of instant messaging (2000) and blogging (2004). CSCW stalwart Bob Kraut has published some of the most influential studies of wikis to date (e.g. Kittur and Kraut, 2008). Even Jonathan Grudin has published widely in the new field including two papers in this year's *International Wiki Symposium*.[2] By no means have MIS and IT been left out completely, but they have not played the leading role one might have predicted *ex ante*.

In a parallel but smaller research stream, a group involved in the broader HCI and CSCW communities has, particularly over the last decade, attempted to explore the antecedents of successful design for creative work. Several high profile contributions to this stream have been the work of Ben Shneiderman. Taking "supporting collaboration" as an explicit component of creativity support (Shneiderman et al., 2006), work on "creativity support tools" has also, even more surprisingly, found itself among a set of researchers who have also historically focused on communities and small groups.

These are broad and highly imperfect generalizations. MIS journals and authors have participated in the research on large group creativity as I have described it. More individualistic approaches to creativity have also been pursued. But if my basic observation is legitimate, we are left asking: Why have small group researchers participated so enthusiastically and effectively in the study of large group creativity?

It is my supposition that because creativity involves recombination of existing materials, it is often best served through connections to others. The Internet has made this possible in increasingly large groups. In that creativity tends to thrive outside of structure and routines that may quash or limit it, creativity tends to looks more like the work of small groups than the work of organizations at the focus of MIS research. Large creative communities, in this sense, represent a type of large "small" group from the perspective of Grudin's typology. CSCW's focus on large group creativity is not an accident. Indeed, is based on this important, if unarticulated, insight and synergy.

Like Grudin's earlier work, this essay is framed around eight challenges for the devel-

[2] http://ur1.ca/0p6uh

Issues around new organizational forms
1. Unclear or "fuzzy" community boundaries
2. Understanding contributor motivations
Issues around collaboration support
3. Providing the "right" amount of communication
4. Building critical mass
Issues around creativity support
5. Personalizing shared work
6. Providing the "right" amount of structure
7. Technological frames constrain as they facilitate
Other issues
8. Evaluation concerns

Table 1: Summary table of the eight challenges for developers of large-scale creative collaboration systems.

opers of systems for large group creativity (LGC). My goal is to act as a guide to researchers and developers working on LGC systems and to map the field in a way that provides support for my suppositions above. I have divided the eight challenges into four groups described in Table 1.

Issues Around New Organizational Forms

The visible ascendancy of LGC, including systems for the production of free, libre, open source software (FLOSS) and projects like Wikipedia, has raised a number of interesting questions for social scientists of organizations and groups. Work at in sociology, economics, and social psychology has asked fundamental questions about how very large computer-mediated communities work, and about how they should be understood and theorized. Key questions around why people join and contribute to these projects remain largely open issues in these literatures. Unsurprisingly, this fundamental ambiguity about what these creative communities are and how they work leads to the most fundamental unresolved challenges to designers' efforts to support them technologically.

Challenge 1: Unclear or "fuzzy" community boundaries. In one of the most influential descriptions of LGC in the social sciences, Benkler (2002, 2006) uses a transaction cost economics (TCE) based approach to the theory of organizations (Coase, 1937; Williamson, 1981) to suggest that communications technology, and the Internet in particular, has driven the cost to contribute to projects like Wikipedia so low that it has necessitated and provoked a novel non-firm, non-market organizational form that Benkler refers to as "peer-production."

Boundaries are a fundamental way to understand the nature of organizations (Santos and Eisenhardt, 2005) and Benkler's TCE-based analysis highlights that LGC communities negotiate their borders with their environments. The question of boundaries has been an issue of concerns in the social science literature on FLOSS more generally (e.g. Ferraro and O'Mahony, 2004; O'Mahony and Bechky, 2008). In their operationalizations of "contributor," researchers have had to decide if a person who fixes a comma on Wikipedia is a user or contributor for the purpose of analysis (e.g. Suh et al., 2009; Ortega, 2009). How about readers (Antin and Cheshire, 2010)? How about a person who vandalizes a page or "trolls" (Shachaf and Hara, 2010)? Community boundaries are, by definition, frustratingly fluid in LGC projects.

These fluid boundaries pose a significant hurdle for designers and a departure from the context of previous work in the design for groups and creativity. Within organizations, informed designers knew (or could know) who they were designing for. For the most part, designers of of LGCs do not. Many of Grudin's (1994b) suggestions for addressing his eight challenges involve customizing for, or educating, groups of users. Similarly, key examples from the literature on creativity support within organizations take clearly defined users for granted (Mamykina et al., 2002; Tiwana and Mclean, 2005).

Researchers have begun to leverage Lave and Wenger's (1991) influential work on communities of practice and legitimate peripheral participation (LPP) to describe multiple types and degrees of membership in a community (e.g. Antin and Cheshire, 2010). However, there remains an important distinction between "peripheral" and "outside" in LPP that poses a challenges in the context of many successful LGCs if researchers are resistant to describing everything as peripheral. Due to this theoretical gap and a paucity of methods for pursuing an answer to these question, researches will continue to struggle to satisfactorily define LGC boundaries.

Challenge 2: Understanding contributor motivations. Given that a community has boundaries, inscrutable and fluid as they may be, a central goal of many LGC designers is to encourage individuals outside of the project's boundaries to move inside – i.e., to motivate, or to increase, participation and contribution. In one of the most influential early papers on FLOSS, Lerner and Tirole (2002) suggest that FLOSS poses largely unanswered economic questions about incentives. Of course, designing to motivate participation requires a model of contributor motivation. Of interest to both theorists and designers, the lack of any such useful model has proved a central challenge to the design of LGC systems.

Survey-based research from FLOSS has pointed to motivations that are both intrinsic and extrinsic, widely divergent, variable both between and within projects, and contingent on a large variety of contextual and environmental factors (e.g. Hars and Ou, 2002; Lakhani and Wolf, 2005). Lakhani and Wolf (2005) and Shah (2006) have suggested, frustratingly for those wishing to encourage volunteer contributions, that many of the most successful FLOSS projects are the products of paid labor. No mystery there.

With such an unclear and unstable foundation, the design literature has also, unsurprisingly, been unfocused and scattered. Ridings and Gefen (2004) reported that motivations to contribute to Internet discussion lists were contingent on the type of discussion. Tedjamulia et al. (2005) describe an eleven-part model of contribution structured around intrinsic and extrinsic motivation but with little empirical support and unclear practical import.

There is wide recognition that we need a better understanding of motivational "levers" and to translate these into designs. But only minor progress has been made toward this goal. That said, recent field experiments like those suggested by Shneiderman (2008) and natural experiments like those by Zhang and Zhu (2010) have made important progress. Zhang and Zhu show that group size affects motivation to contribute to Chinese Wikipedia. Technical changes to online communities like Community Lab's *MovieLens* (e.g. Beenen et al., 2004) provide first steps in that group's explorations of the role that emphasizing different aspects of contributions and goal-setting play in motivating participation.

Issues Around Support of Collaboration

Work in CSCW over the last two decades have produced a large body of knowledge on the computer support of group collaborative activity. That said, a move from the traditional context of CSCW – small groups – to large collaborations complicates, and may even render irrelevant or reverse, establish findings. Additionally, the problem of supporting creativity can conflict with the goal of supporting collaboration. Collaboration in LGCs tends to be largely unstructured and to work like small groups or teams. But it is also likely to differ in important ways that remain poorly understood.

Challenge 3: Providing the "right" amount of communication. Work in CSCW has focused heavily on support for communication. Classics in CSCW including Hutchins (1990) on team navigation, Clark and Brennan (1992) on grounding, and Tatar et al. (1991) on the *Cognoter* project at Xerox PARC, each emphasized the positive role that increased and improved communication could play in facilitating group work. The awareness substream built on these fundamental insight has suggested that information on group activity can help provide context and lead to smoother and more effective collaborative work in creative contexts (Dourish and Bellotti, 1992; Gutwin and Greenberg, 2002).

That said, the most naive awareness models show brittleness as we move them from the small teams typical in classic CSCW examples – *Cognoter* was a brainstorming tool built for 3 users – to the huge groups in LGC projects like Wikipedia. We have every reason to believe that inter-collaborator communication still plays a critical role in LGCs, but the overwhelming size of some collaborations means we must be more selective. Gutwin and Greenberg (1998) has described a tension around awareness between individual and group support; the problem they describe is both more widespread and more acute in LGCs.

Viegas et al. (2004) has shown that providing awareness on activity in projects like Wikipedia is possible, but that it requires careful selection of data to display and creative visualization. But while Viégas et al. point to one solution for Wikipedia editors, we lack generalizable knowledge of how to design awareness interfaces for LGC. Clearly, the classic "more communication is better" models from CSCW will result in information overload

(Hiltz and Turoff, 1985) but we lack theory, or even heuristics, to strike the right balance between too much and too little information.

Challenge 4: Building critical mass. In his second major challenge, Grudin (1994b) suggested that building critical mass was a problem for groupware systems in the more traditional organizational context. This continues to be a problem in the context of LGC systems, although the nature of the problem is different. Solutions are also rendered both dramatically different and are likely to be more difficult.

The basic problem of critical mass certainly seems relevant to LGCs: Healy and Schussman (2003) have shown that less than half of all FLOSS projects hosted on SourceForge attract a second contributor. More than 95% of wikis on Wikia fail to achieve critical mass (Kittur and Kraut, 2010). Data from Scratch shows that only ten percent of projects are remixed after a year on the site. Grudin's solutions to critical mass problems involved top-down managerial incentives for the use of groupware systems. Volunteer driven peer-production projects clearly require a different and more nuanced approach.

Theory from social movements on resource mobilization (McCarthy and Zald, 1977) and on framing (Snow et al., 1986; Benford and Snow, 2000) provide theoretical leaping-off points. Work from political science on collective action (Olson, 1965; Oliver et al., 1985; Oliver and Marwell, 1988) provides additional context. That said, translating insights from these literatures into designs remains an largely unaddressed challenge.

We must first build empirical knowledge of how and why some LGC projects attain critical mass while most do not. A good technique – still largely unexplored in the broader literature on LGCs – involves building samples of both successful and unsuccessful projects. Most research has focused on only the ones that work. Designers need variation in critical mass attainment if they want to understand why some LGCs achieve it while most do not.

Issues Around Support of Creativity

The design literature on creativity is both younger and less well developed than the design literature on collaboration. Aggravating (and perhaps influencing) this fact, the basic concept of creativity is more slippery and more difficult to pin down. As a result, we know less about how to support creativity than we do about how to support collaboration. To complicate matters, the knowledge we do have about creativity support can lead to prescriptions that are at odds with concepts from the literature on the support of collaboration. LGCs must support both creativity and collaboration. Understanding how to do both, without sacrificing either, reveals a major set of challenges.

Challenge 5: Technological frames constrain as they facilitate. Orlikowski (1992) and Orlikowski and Gash (1994) each build on Goffman's (1974) work on framing to suggest that cognitive "technological frames" can limit a groups' up-take of a system. In Orlikowski's empirical example, users within a large consulting firm adopted and began to use the *Lotus Notes* groupware systems but lacked a frame associated with a collaborative knowledge sharing system. As a result, they used *Notes* exactly as they had used previous email systems without groupware functionality. They simply slotted *Notes* into their existing technological frame and acted accordingly. Frames serve to reduce users' cognitive load by constraining one's imagination of one might conceivably do with a system. If a user only has an "email" frame, she will use "more-than-email" systems as if they were less capable tools.

Support for creativity is, by definition, about building the support for multiple possible use cases. For example, Shneiderman et al. (2006) suggests facilitating multiple paths in their third principle on support of creative thinking. Shneiderman (2002) suggest support for new associations and explorations in his eight tasks that creativity tools should support. This sets up an important tension with the role that frames play in getting users on the same page – a process at the heart of coordination (Malone and Crowston, 1990; Crowston, 1997) and collaboration by extension.

Wikipedia may have been easier for users to contribute to because it leveraged a centuryold frame associated with encyclopedia writing (Reagle, 2008, 2009) that helped get users on the same page with each other and to coordinate. That said, Wikipedia's growth has stagnated (Suh et al., 2009; Ortega, 2009) and most new articles are deleted because they do not match the encyclopedia frame and are deemed "non-notable" and "unencyclopedic." The frame which best facilitates coordination and collective action will often be poorly suited to the most creative users. Designers have yet to address this tension in any satisfactory depth and it remains an important area for further research.

Challenge 6: Providing the "right" amount of structure. The previous challenge in regards to framing is about striking the balance between too much and too little *cognitive* structure. There is evidence that a similar balance should exist between too much and too little *technological* structure. Using ethnographic methods, Suchman (1983) describes the highly contingent and complicated nature of office work, in extreme depth, as a way of illustrating the difficulties of supporting workflow.

Suchman (1994) is highly critical of Winograd's (1986) "speech act theory" approach to modeling group behavior as overly static, limiting, and dangerous in the hands of management who would not understand or respect its limitations. Convincing or not in its defense of speech act theory, (Winograd, 1994) offers a nuanced response to Suchman that captures the complex balance that designers should attempt to strike between building technological support for particular types of work – a process inherent in the design of any technological system – and the often invisible social and political problems associated with technological structure that is, or becomes, too constraining.

Just as creative work may need more cognitive flexibility, technological support for creative work may have deleterious effect on creativity. Work on bricolage in organizations has suggested that the best performing organizations practice creative bricolage, but use it in moderation (Miner et al., 2001). Sawyer (2006) emphasizes the idea of "structured improvisation" in support of increased creativity in education. Simulation work on structure in organizations also suggests a "happy medium" and that it may be safer to err toward too much structure (Davis et al., 2009).

But this common calculus is confused when the performance metric is, or includes, creativity. Intuition suggests that designers might want to err away from structure in these cases. Successful LGC projects like USENET offer little in the way of structured technological support and provide extraordinarily flexible system for work (Kollock and Smith, 1996; Whittaker et al., 1998). We lack both theory and empirical evidence critical to striking this balance optimally.

Challenge 7: Personalizing shared work. Work on creativity and learning has emphasized the importance of personally meaningful tasks in deep creative engagement (Csikszentmihalyi, 1990; Papert, 1980; Jonassen, 1999). Once again, this sets up a tension with the literature on collaborative work. In several core definitions, the literature on CSCW implies shared work products and goals as sites for collaboration (Ellis et al., 1991; Kling, 1991) or as prerequisites to coordination (Malone and Crowston, 1990). If a creative product is extremely personally meaningful, it is likely (although clearly not necessarily) going to be less interesting to others. A suit that is perfectly tailored to one person is less likely to fit her friend at all and is almost guaranteed to fit less well.

Several approaches may mitigate this concern. von Hippel and Katz (2002) have suggested "toolkits" as a model that involves the use of sets of common components, pieces, or tools. Where hardware imposes the need for at least some commonality and resuability between electronic objects, this approach has been used effectively in the support of creativity (e.g. Greenberg and Fitchett, 2001; Buechley et al., 2008). In software based creative systems, users have suggested a related model of a knowledge repository useful as a library for a variety of tasks (e.g. Scardamalia and Bereiter, 1993) or through highly modular design that allows users to "own" and personalize parts of a larger integrated system (e.g. Bruckman, 1998).

Each approach offers trade-offs between the level, degree, and type of personalization possible, and the degree and nature of collaborative production available to users of the system. Understanding these trade-offs, and the settings and contexts where we might want to pursue one approach over another, remains an important challenge for designers of LGCs.

Other issues

Challenge 8: Evaluation concerns. Included on Grudin's 1994 list, evaluation remains as critical a problem in LGC systems as it was for groupware. Like groupware, LGC systems should be more meaningfully evaluated in terms of their actual use by users, "in the wild." In their report on creativity support tools, Shneiderman et al. (2006) emphasized difficulties of evaluation as well.

That said, success for an LGC is different than for a groupware system. At a basic level, the ability of a project to attract contributors and participants provides a meaningful benchmark that must be interpreted differently for groupware where use was often mandated by management. Evaluating the outputs of LGCs is at least as intractable. In general, LGCs include all the of the difficulties of evaluation associated with CSCW and all of the difficulties associated with support of creativity. Additional interaction effects are likely to complicated the situation even further.

Shneiderman et al. (2006) point to longitudinal analysis and field experiments as useful methods of addressing problems of evaluation for creativity support tools which seem relevant to LCGs more generally. This represents good advice about evaluation methods but seems to have not been followed up much in practice. Designers have paid lip service to the problem and the idea of more sophisticated evaluation methods, but continue to evaluate their systems with usability studies because they are readily available, relatively cheap when compared to longitudinal methods and field experiments, and provide clear path to publication. As a result, we have made very little progress since Grudin's statement nearly twenty years ago.

Conclusion

Many of the problems discussed in this essay can be seen as generalizations of the specific problem highlighted by Gutwin and Greenberg (1998) in the tension between supporting individuals and supporting groups. Creativity, in most approaches, remains a highly individualized concept at the moment of creation. We understand that creativity involves the recombination of materials and that networks and collaboration can help the process along, but the creative act itself is still attributed, in almost all cases, to an individual. A parallel issue is associated with learning. We may learn in and through groups, but learning still happens within an individual. One solution to this conundrum would be to think of organizational creativity in the same sense we think of organizational learning (Levitt and March, 1988). Indeed, the literature on innovation in organizations suggests several potentially meaningful approaches in this regard that we might explore through analogy. Of course, organizational translation of these concepts into the realm of technology support will still present many of the challenges described above.

DESIGN PRINCIPLES FOR GROUP CREATIVITY

You have been hired as a consultant to a startup that is developing "tools to support group creativity." You are asked to write a memo describing "guiding principles" to guide the design of the company's technologies, drawing on research on collaboration and creativity. What would you write?

Introduction

This essay is an attempt to distill a set of high-level design principles from the scholarly literature on technological design for collaboration and creativity. It is by no means the first attempt to do so. Noteworthy similar attempts include Nickerson (1999) who described twelve steps to teaching creativity, Shneiderman (2002) who described eights type of collaborative support, and Shneiderman et al. (2006) who offered twelve general principles for the support of creativity. That said, this list has different goals than its predecessors. First, this list attempts to build off this previous work. Second, it tries to offer more specific advice that previous attempts which, for example, suggest designers "build motivation" and "support exploration." Finally, it attempts to bring an added focus to issues of collaborative creativity where previous attempts (e.g., Shneiderman et al.) suggest that designers, "support collaboration;" as I have shown in the preceding essay, the addition of collaboration can complicate the support of creativity.

A list of eight guiding design principles for tools supporting group creativity is provided to complement the list of eight challenges in the preceding section. That said, no effort was made to structure a correspondence between challenges and principles in the two texts. A list of guiding principles is distilled into Table 2. Although each of the principles is phrased in non-academic terms that should be accessible to a manager at my hypothetical start-up, each principle stems from research in the academic study of collaboration, creativity, and innovation, and organizations. Citations in the descriptions of each should speak to these principles' provenance.

1. Your users understand their problem better than you do.
2. There should be more than one way to do it.
3. Supporting sociality in service of creativity.
4. Break projects down into digestible pieces.
5. Focus on working outputs, not toys.
6. Provide context and history.
7. Avoid walled gardens.
8. All of this applies to you too.

Table 2: List of eight guiding design principles for systems supporting group creativity.

Principle 1. Your users understand their problem better than you do.

Most definitions of creativity suggest the a creative act is a recombination of material in a novel way. Creative actions are responses to particular problems, curiosities, needs, or desires of the creator. Designers of tools to support creative work build technologies to let their *users* make recombinations more easily themselves.

Traditional product design methodology emphasizes intensive processes of market research in order to understand users' needs (Ulrich and Eppinger, 1994). No researcher has done a better of showing how much better users understand their own problems than von Hippel (2005). A better goal than trying to understand users' needs to is build tools for users so that they can solve their own problems. This is the power – and the point – of creativity support tools. This should be the central guiding principle for any designer.

Minor assumptions about how a user want to act or what they want to create will constrain outputs significantly and will likely restrict use of a tool to certain groups and certain types of use. Buechley and her colleague's (2008; 2010) work on the LilyPad platform shows how "superficial" changes to a microcontroller have opened up the tool to a wide variety of new uses by women who were historically excluded from creative work with those tools. von Hippel and Katz's (2002) work on toolkits, of which Buechley's work is influenced, provides and important example of what design for user-directed creativity can look like and what it can accomplish.

Principle 2. There should be more than one way to do it.

Shneiderman et al.'s (2006) third principle suggests that designers should, "support many paths and many styles." In the collaborative context, this means *simultaneous* support for these divergent, forking, or just different paths. On the one hand, this means supporting multiple types of users. This is strongly emphasized by Papert (1980) and by Gardner (1993) who goes further to suggest that are different types of creativity we should aim to support. Support for creativity, by definition, means we cannot know what will be produced. With details up in the air, it is easy to impose constraints that limit possible processes and outputs. Principle 2 reminds us to stay flexible.

Even one particular user is likely to grown and change in their patterns of use of a system over time in a collaborative context and tool will simultaneously need to host users at different stages or points of their practice. Lave and Wenger (1991) provides a useful description of communities with multiple levels and types of engagement as well as an apprenticeship inspired model for how individuals in communities of practice can move from peripheral engagement to more core forms of contribution. Building on Lave and Wenger, Antin and Cheshire (2010) argue that readers play an important role in Wikipedia as a way of dispelling the widespread notion that readers are "free riding" on Wikipedia editors work. There are multiple ways to contribute to most creative products and some of them may not look like contributions to us at all.

Fish et al.'s (1988) work on the *Quilt* collaborative editing system, and a variety of other systems supporting collaborative writing, have been criticized as relying too heavily on workflow for codifying involvement through roles and tasks-based structure. Ellis et al. (1991) and others convincingly argued that the most successful collaborative writing systems allow for less structured collaboration and a variety of types of participation styles that, when implemented successfully, allow users to make multiple types of contributions over time and to interact in novel and unanticipated ways.

Principle 3. Support sociality in service of creativity.

A large body of work in CSCW has emphasized the importance of communication in collaboration (e.g Tatar et al., 1991). Work building on this fundamental insight has emphasized the related construct of awareness (Dourish and Bellotti, 1992; Gutwin and Greenberg, 2002). Erickson and Kellogg's (2000) concept of social translucence provides another iteration on the same basic concept. Increased communication and information exchange both presents additional information necessary to coordination (Malone and Crowston, 1990; Crowston, 1997) and the raw material for creative recombination.

The most successful and enduring tools in collaborative creativity are often the most social. Finholt and Sproull's (1990) work on email groups in organizations, Nardi et al.'s (2004) work on blogs, Boyd and Ellison's (2007) work on social networks, and Nardi et al.'s (2000) work on instant messaging have all emphasized the social component of these systems. They are tools that are often described purely in terms of their use for socialization. That said, although their success have led us to take them for granted, this is a list of the best technological tools available for the supporting of creativity.

Of course, socialization can come at the price of productivity in more traditional terms. In the context of creativity, this trade-off may be pushed in the direction of more sociality. Social interaction provides both the means and materials for creativity through the creation of the tightly joined networks through which to transfer resources necessary for creation (e.g. Obstfeld, 2005) and coordinate work on creative task execution. Although an instant messaging system is likely to be very social in use and although it will provide less specific support for explicit creative production, the addition of functionality for social interaction through instant messaging is a likely piece of many successful group creativity systems precisely because of its capacity for sociality.

Principle 4. Break projects down into digestible pieces.

A variety of research perspectives suggests that breaking down larger projects into manageable chunks is an important piece of support for group creativity. On the one hand, Papert (1980) argues strongly for, and Abelson and diSessa (1981) implements, a model of creativity and learning with the capacity for enormous complexity that is slowly built up out of small "mind-size bites" and basic, comprehensible propositions. Abelson and diSessa implement a system to do general relativity calculations in what some might dismiss as a child's programming language with few routines of more than a dozen lines and no concept more complicated than recursion.

In the management and software engineering literature, MacCormack et al. (2006, 2008) have used Baldwin and Clark's (2000) concept of modularity to show that FLOSS projects that are the sites of more active collaborative creation tend to be much more modular than functionally similar closed equivalents. Although the exact causal mechanism is not possible to determine in their design, MacCormack et al. show that work in what are effectively cognitively smaller chunks of code are more common in more collaborative FLOSS projects than in their proprietary equivalents.

Work from CSCW also suggests that successful group creativity is likely to happen in projects which can be broken down by users into manageable chunks and dealt with in a less structured fashion (Ellis et al., 1991). I have shown in my own work on the Scratch online community (Resnick et al., 2009) that users are less receptive to collaborative interaction in social media sharing environments when the shared chunks are larger (Hill et al., 2010).

Principle 5. Focus on working outputs, not toys.

Papert (1980) strongly cautious against the lure of "toy" versions that might serve to "dumb down" powerful ideas as (false) trade-offs for accessibility. Better, Papert suggests, is for designers to support learning and creativity with working "real" versions that can scale up and grow – through the assemblage of smaller pieces described in the previous section – into larger more complicated outputs that ultimately become the final working creative products. A variety of other perspective support the idea of this model for innovation, creativity and cooperation.

In the literature on product development, Wheelwright and Clark (1992) have argued strongly in favor of more and faster prototyping at the center of the product development process. They argue that by facilitating the creation of working interactions more quickly and more often, designers are able to debug earlier and get their product working and onto the market faster – an observation supported in the product development literature more generally (Ulrich and Eppinger, 1994).

Early stage working outputs, often in the form of prototypes, can facilitate increased interaction (a principle suggested by Shneiderman et al.) around concrete examples. These systems have also been effectively used in a variety of brainstorming environments (e.g. Sutton and Hargadon, 1996). These working iterations can create common ground and context (see the next section) and ensure that important disagreements are not hidden by convenient abstractions.

Principle 6. Provide context and history.

Suchman's (1995) article on "making work visible" offers a description of how technological support of work, and the studies of work as an antecedent, offer limited abstractions of activities that end up constraining, rather than facilitating, the nature of the work being done. Provision of context and history – essentially, efforts to make the representation more like the "final" product – are an important means of mitigating this ultimately unavoidable issue.

The literature on awareness describes on important means of providing context. History is another important aspect of work that can provide another important type. Hill et al. (1992) use the metaphor of "wear" – imagine the wear on a book or recipe card that helps one find the page they are looking for – as a way to think about the design of technological systems by making history more visible to users and facilitating collaboration more effectively.

In creative work, history also provides additional raw materials from within the system for creative recombination. In collaborative work, history provides context and a means for improved coordination through increased "social translucency" (Erickson and Kellogg, 2000). Tools for group creativity are likely to benefit from both.

Principle 7. Avoid walled gardens.

If, as the previous principle suggests, history and context internal to a collaborative interaction or a technological system can improve group creativity by providing more raw material for recombination, additional context from "outside" can provide an even richer and more diverse array of resources.

Empirical research on remixing confirms the creative potential of a wide variety of sources including access to those outside of any particular community or system. Diakopoulos et al. (2007) show that the group video re-remixing tool *Jumpcut* succeeded through its ability to bring in external content. Relatively limited remixing of content on the San Francisco Film Festival website might be due in part to the fact most of the content people want to remix remains outside of, and inaccessible from, the SIFF remixing tool (Shaw and Schmitz, 2006).

Shneiderman (2002) describes "collection" as one of the four major steps in the creative process and urges large assemblages of materials within creativity support tools. Building more isolated "walled gardens" makes consistency and integration of content easier but it is likely to introduce limits that will be unacceptably stringent for all but the most constrained creative processes.

Principle 8. All of this applies to you too.

Shneiderman et al. (2006) urge designers to "design for designers" and to "evaluate your tools." An synthetic and reflective version of these imperatives suggests that designers would benefit by applying these principles to the design of their own systems. The group creativity tool is – or can be – a toolkit in von Hippel's (2005) sense. The most powerful design tools – in Papert's terms, at least – are systems that can design design tools. As Suchman (1995) reminds us, any tool will constrain creativity and present limits to the nature of collaborative and creative production. However, to the degree that the tool itself is modifiable and can became not just the site of, but the raw materials for, creative recombination, the system can exceed the creative capacities of even the most creative designer. Working together, all your users are more creative than you.

REFERENCES

Abelson, H. and A. diSessa (1981, June). *Turtle geometry: The computer as a medium for exploring*. The MIT Press.

Antin, J. and C. Cheshire (2010). Readers are not free-riders: reading as a form of participation on wikipedia. In *Proceedings of the 2010 ACM conference on Computer supported cooperative work*, Savannah, Georgia, USA, pp. 127–130. ACM.

Baldwin, C. Y. and K. B. Clark (2000, March). *Design Rules, Vol. 1: The Power of Modularity*. The MIT Press.

Beenen, G., K. Ling, X. Wang, K. Chang, D. Frankowski, P. Resnick, and R. E. Kraut (2004). Using social psychology to motivate contributions to online communities. In *Proceedings of the 2004 ACM conference on Computer Supported Cooperative Work*, Chicago, Illinois, USA, pp. 212–221. ACM.

Benford, R. D. and D. A. Snow (2000). Framing processes and social movements: An overview and assessment. *Annual Review of Sociology 26*(1), 611–639.

Benkler, Y. (2002). Coase's penguin, or, linux and the nature of the firm. *Yale Law Journal 112*(3), 369.

Benkler, Y. (2006, May). *The Wealth of Networks: How Social Production Transforms Markets and Freedom*. Yale University Press.

Boyd, D. M. and N. B. Ellison (2007, November). Social network sites: Definition, history, and scholarship. *Journal of Computer-Mediated Communication 13*(1), 210–230.

Bruckman, A. (1998, March). Community support for constructionist learning. *Computer Supported Cooperative Work (CSCW) 7*(1), 47–86.

Buechley, L., M. Eisenberg, J. Catchen, and A. Crockett (2008). The LilyPad arduino: using computational textiles to investigate engagement, aesthetics, and diversity in computer science education. In *Proceeding of the twenty-sixth annual SIGCHI conference on Human factors in computing systems*, Florence, Italy, pp. 423–432. ACM.

Buechley, L. and B. M. Hill (2010). LilyPad in the wild. Working Paper.

Clark, H. H. and S. E. Brennan (1992, December). Grounding in communication. In R. M. Baecker (Ed.), *Readings in Groupware and Computer-Supported Cooperative Work: Assisting Human-Human Collaboration* (1 ed.). Morgan Kaufmann.

Coase, R. H. (1937, November). The nature of the firm. *Economica 4*(16), 386–405.

Crowston, K. (1997, April). A coordination theory approach to organizational process design. *Organization Science 8*(2), 157–175.

Csikszentmihalyi, M. (1990, March). *Flow the Psychology of Optimal Experience* (1st ed.). Harpercollins Publisher.

Davis, J. P., K. M. Eisenhardt, and C. B. Bingham (2009). Optimal structure, market dynamism, and the strategy of simple rules. *Administrative Science Quarterly 54*, 413–452.

Diakopoulos, N., K. Luther, Y. E. Medynskiy, and I. Essa (2007). The evolution of authorship in a remix society. In *Proceedings of the eighteenth conference on Hypertext and hypermedia*, Manchester, UK, pp. 133–136. ACM.

Dourish, P. and V. Bellotti (1992). Awareness and coordination in shared workspaces. In *CSCW '92: Proceedings of the 1992 ACM conference on Computer-supported cooperative work*, New York, NY, USA, pp. 107–114. ACM Press.

Ellis, C. A., S. J. Gibbs, and G. Rein (1991). Groupware: Some issues and experiences. *Communications of the ACM 34*(1), 39–58.

Erickson, T. and W. A. Kellogg (2000). Social translucence: an approach to designing systems that support social processes. *ACM Trans. Comput.-Hum. Interact. 7*(1), 59–83.

Ferraro, F. and O'Mahony (2004, March). Managing the boundary of an 'Open' project. Harvard NOM Working Paper 03-60, Harvard University.

Finholt, T. and L. S. Sproull (1990). Electronic groups at work. *Organization Science 1*(1), 41–64.

Fish, R. S., R. E. Kraut, and M. D. P. Leland (1988). *Quilt: a collaborative tool for cooperative writing*. Palo Alto, California, United States: ACM Press.

Gardner, H. E. (1993, July). *Creating Minds: An Anatomy Of Creativity Seen Through The Lives Of Freud, Einstein, Picasso, Stravinksy, Eliot, Graham, and Gandhi*. Basic Books.

Goffman, E. (1974, December). *Frame Analysis: An Essay on the Organization of Experience* (First Edition ed.). Harvard Univ Pr.

Greenberg, S. and C. Fitchett (2001). Phidgets: easy development of physical interfaces through physical widgets. In *Proceedings of the 14th annual ACM symposium on User interface software and technology*, Orlando, Florida, pp. 209–218. ACM.

Grudin, J. (1994a). Computer-supported cooperative work: History and focus. *IEEE Computer 27*(5), 19–26.

Grudin, J. (1994b). Groupware and social dynamics: Eight challenges for developers. *Commun. ACM 37*(1), 92–105.

Gutwin, C. and S. Greenberg (1998). Design for individuals, design for groups: Tradeoffs between power and workspace awareness. In *Proceedings of the 1998 ACM conference on Computer supported cooperative work*, Seattle, Washington, United States, pp. 207–216. ACM.

Gutwin, C. and S. Greenberg (2002). A descriptive framework of workspace awareness for Real-Time groupware. *Computer Supported Cooperative Work (CSCW) 11*(3), 411–446.

Hars, A. and S. Ou (2002). Working for free? motivations for participating in open-source projects. *International Journal of Electronic Commerce 6*(3), 25–39.

Healy, K. and A. Schussman (2003). The ecology of Open-Source software development.

Hill, B. M., A. Monroy-Hernández, and K. Olson (2010). Responses to remixing on a social media sharing website. In *Proceedings of the Fourth International AAAI Conference on Weblogs and Social Media*, Washington, D.C. AAAI.

Hill, W. C., J. D. Hollan, D. Wroblewski, and T. McCandless (1992). Edit wear and read wear. In *CHI '92: Proceedings of the SIGCHI conference on Human factors in computing systems*, New York, NY, USA, pp. 3–9. ACM Press.

Hiltz, S. R. and M. Turoff (1985). Structuring computer-mediated communication systems to avoid information overload. *Commun. ACM 28*(7), 680–689.

Hutchins, E. (1990, May). The technology of team navigation. In J. Galegher, R. E. Kraut, and C. Egido (Eds.), *Intellectual Teamwork: Social and Technological Foundations of Cooperative Work* (1 ed.). Psychology Press.

Jonassen, D. H. (1999). Designing constructivist learning environments. In *Instructional design theories and models: A new paradigm of instructional theory*, pp. 371–396.

Kittur, A. and R. E. Kraut (2008). Harnessing the wisdom of crowds in wikipedia: quality through coordination. San Diego, CA, USA, pp. 37–46. ACM.

Kittur, A. and R. E. Kraut (2010). Beyond wikipedia: coordination and conflict in online production groups. In *Proceedings of the 2010 ACM conference on Computer supported cooperative work*, Savannah, Georgia, USA, pp. 215–224. ACM.

Kling, R. (1991). Cooperation, coordination and control in computer-supported work. *Commun. ACM 34*(12), 83–88.

Kollock, P. and M. Smith (1996). Managing the virtual commons: Cooperation and conflict in computer communities. *Computer-Mediated Communication: Linguistic, Social, and Cross-Cultural Perspectives*, 109–128.

Lakhani, K. and B. Wolf (2005). Why hackers do what they do: Understanding motivation and effort in Free/Open source software projects. In J. Feller, B. Fitzgerald, S. A. Hissam, and K. R. Lakhani (Eds.), *Perspectives on Free and Open Source Software*, pp. 3–22. MIT Press.

Lave, J. and E. Wenger (1991, September). *Situated Learning: Legitimate Peripheral Participation*. Cambridge University Press.

Lerner, J. and J. Tirole (2002). Some simple economics of open source. *Journal of Industrial Economics 50*(2), 197–234.

Levitt, B. and J. G. March (1988). Organizational learning. *Annual Review of Sociology 14*(1), 319–338.

MacCormack, A., J. Rusnak, and C. Y. Baldwin (2006). Exploring the structure of complex software designs: An empirical study of open source and proprietary code. *Management Science 52*(7), 1015–1030.

MacCormack, A. D., J. Rusnak, and C. Y. Baldwin (2008, March). Exploring the duality between product and organizational architectures: A test of the mirroring hypothesis. working paper 08-039, Harvard Business School.

Malone, T. W. and K. Crowston (1990). What is coordination theory and how can it help design cooperative work systems? In *Proceedings of the 1990 ACM conference on Computer-supported cooperative work*, Los Angeles, California, United States, pp. 357–370. ACM.

Mamykina, L., L. Candy, and E. Edmonds (2002). Collaborative creativity. *Communications of the ACM 45*(10), 96–99.

McCarthy, J. D. and M. N. Zald (1977, May). Resource mobilization and social movements: A partial theory. *The American Journal of Sociology 82*(6), 1212–1241. ArticleType: primary_article / Full publication date: May, 1977 / Copyright © 1977 The University of Chicago Press.

Miner, A. S., P. Bassoff, and C. Moorman (2001, June). Organizational improvisation and learning: A field study. *Administrative Science Quarterly 46*(2), 304–337.

Nardi, B. A., D. J. Schiano, and M. Gumbrecht (2004). Blogging as social activity, or, would you let 900 million people read your diary? In *Proceedings of the 2004 ACM conference on Computer supported cooperative work*, Chicago, Illinois, USA, pp. 222–231. ACM.

Nardi, B. A., S. Whittaker, and E. Bradner (2000). Interaction and outeraction: Instant messaging in action. In *Proceedings of the 2000 ACM conference on Computer supported cooperative work*, Philadelphia, Pennsylvania, United States, pp. 79–88. ACM.

Nickerson, R. (1999). Enhancing creativity. In R. J. Sternberg (Ed.), *Handbook of Creativity* (1 ed.)., pp. 392–430. Cambridge University Press.

Obstfeld, D. (2005, March). Social networks, the tertius iungens orientation, and involvement in innovation. *Administrative Science Quarterly 50*(1), 100–130.

Oliver, P., G. Marwell, and R. Teixeira (1985, November). A theory of the critical mass. i. interdependence, group heterogeneity, and the production of collective action. *The American Journal of Sociology 91*(3), 522–556.

Oliver, P. E. and G. Marwell (1988, February). The paradox of group size in collective action: A theory of the critical mass. II. *American Sociological Review 53*(1), 1–8.

Olson, M. (1965, January). *The Logic of Collective Action: Public Goods and the Theory of Groups* (First Edition ed.). Harvard University Press.

O'Mahony, S. and B. A. Bechky (2008, September). Boundary organizations: Enabling collaboration among unexpected allies. *Administrative Science Quarterly 53*(3), 422–459. Special Issue: Social Movements in Organizations and Markets.

Orlikowski, W. J. (1992). Learning from notes: organizational issues in groupware implementation. In *Proceedings of the 1992 ACM conference on Computer-supported cooperative work*, Toronto, Ontario, Canada, pp. 362–369. ACM.

Orlikowski, W. J. and D. C. Gash (1994). Technological frames: making sense of information technology in organizations. *ACM Transactions of Information Systems 12*(2), 174–207.

Ortega, F. (2009, April). *Wikipedia: A Quantitative Analysis*. PhD dissertation, Universidad Rey Juan Carlos.

Papert, S. (1980). *Mindstorms : children, computers, and powerful ideas*. New York: Basic Books.

Reagle, J. (2009). Wikipedia: the happy accident. *interactions 16*(3), 42–45.

Reagle, J. M. (2008). *In Good Faith: Wikipedia Collaboration and the Pursuit of the Universal Encyclopedia*. PhD, New York University.

Resnick, M., B. Silverman, Y. Kafai, J. Maloney, A. Monroy-Hernández, N. Rusk, E. Eastmond, K. Brennan, A. Millner, E. Rosenbaum, and J. Silver (2009). Scratch: programming for all. *Communications of the ACM 52*(11), 60.

Ridings, C. M. and D. Gefen (2004). Virtual community attraction: Why people hang out online. *Journal of Computer-Mediated Communication 10*(1), 00–00.

Santos, F. M. and K. M. Eisenhardt (2005, September). Organizational boundaries and theories of organization. *Organization Science 16*(5), 491–508.

Sawyer, R. K. (2006, April). Educating for innovation. *Thinking Skills and Creativity 1*(1), 41–48.

Scardamalia, M. and C. Bereiter (1993). Computer support for knowledge-building communities. *The Journal of the Learning Sciences 3*(3), 265–283.

Shachaf, P. and N. Hara (2010). Beyond vandalism: Wikipedia trolls. *J. Inf. Sci. 36*(3), 357–370.

Shah, S. K. (2006, July). Motivation, governance, and the viability of hybrid forms in open source software development. *Management Science 52*(7), 1000–1014.

Shaw, R. and P. Schmitz (2006). Community annotation and remix: a research platform and pilot deployment. In *Proceedings of the 1st ACM international workshop on Human-centered multimedia*, Santa Barbara, California, USA, pp. 89–98. ACM.

Shneiderman, B. (2002). Creativity support tools. *Communications of the ACM 45*(10), 116–120.

Shneiderman, B. (2007). Web science: a provocative invitation to computer science. *Commun. ACM 50*(6), 25–27.

Shneiderman, B. (2008, March). Science 2.0. *Science 319*(5868), 1349–1350.

Shneiderman, B., G. Fischer, M. Czerwinski, M. Resnick, B. Myers, L. Candy, E. Edmonds, M. Eisenberg, E. Giaccardi, H. Hewett, P. Jennings, B. Kules, K. Nakakoji, J. Nunamaker, R. Pausch, T. Selker, E. Sylvan, and M. Terry (2006). Creativity support tools: Report from a US national science foundation sponsored workshop. *International Journal of Human-Computer StInteractionudies 20*(2), 61–77.

Snow, D. A., E. B. Rochford, S. K. Worden, and R. D. Benford (1986, August). Frame alignment processes, micromobilization, and movement participation. *American Sociological Review 51*(4), 464–481.

Suchman, L. (1994). Do categories have politics? *Computer Supported Cooperative Work (CSCW) 2*(3), 177–190.

Suchman, L. (1995). Making work visible. *Communications of the ACM 38*(9), 56–ff.

Suchman, L. A. (1983). Office procedure as practical action: Models of work and system design. *ACM Trans. Inf. Syst. 1*(4), 320–328.

Suh, B., G. Convertino, E. H. Chi, and P. Pirolli (2009). The singularity is not near: slowing growth of wikipedia. In *Proceedings of the 5th International Symposium on Wikis and Open Collaboration*, Orlando, Florida, pp. 1–10. ACM.

Sutton, R. I. and A. Hargadon (1996, December). Brainstorming groups in context: Effectiveness in a product design firm. *Administrative Science Quarterly 41*(4), 685–718.

Tatar, D. G., G. Foster, and D. G. Bobrow (1991). Design for conversation: Lessons from cognoter. In *Computer-supported cooperative work and groupware*, pp. 55–80. Academic Press Ltd.

Tedjamulia, S. J. J., D. L. Dean, D. R. Olsen, and C. C. Albrecht (2005). Motivating content contributions to online communities: Toward a more comprehensive theory.

In *System Sciences, 2005. HICSS'05. Proceedings of the 38th Annual Hawaii International Conference on*, pp. 193b–193b.

Tiwana, A. and E. R. Mclean (2005). Expertise integration and creativity in information systems development. *J. Manage. Inf. Syst. 22*(1), 13–43.

Ulrich, K. T. and S. D. Eppinger (1994). *Product Design and Development*. New York NY: McGraw-Hill.

Viegas, F. B., M. Wattenberg, and K. Dave (2004). Studying cooperation and conflict between authors with history flow visualizations. In *CHI '04: Proceedings of the SIGCHI conference on Human factors in computing systems*, New York, NY, USA, pp. 575–582. ACM Press.

von Hippel, E. (2005). *Democratizing Innovation*. Cambridge, Massachusetts: The MIT Press.

von Hippel, E. and R. Katz (2002, July). Shifting innovation to users via toolkits. *Management Science 48*(7), 821–833.

Weick, K. E. (1995, May). *Sensemaking in Organizations* (1 ed.). Sage Publications, Inc.

Wheelwright, S. C. and K. Clark (1992, June). *Revolutionizing Product Development: Quantum Leaps in Speed, Efficiency, and Quality*. Free Press.

Whittaker, S., L. Terveen, W. Hill, and L. Cherny (1998). The dynamics of mass interaction. In *Proceedings of the 1998 ACM conference on Computer supported cooperative work*, Seattle, Washington, United States, pp. 257–264. ACM.

Williamson, O. E. (1981, November). The economics of organization: The transaction cost approach. *The American Journal of Sociology 87*(3), 548–577.

Winograd, T. (1986). *Understanding Computers and Cognition : A New Foundation for Design*. Greenwood Publishing Group, Incorporated.

Winograd, T. (1994). Categories, disciplines, and social coordination. *Computer Supported Cooperative Work (CSCW) 2*(3), 191–197.

Zhang, X. M. and F. Zhu (2010). Group size and incentives to contribute: A natural experiment at chinese wikipedia. *American Economic Review*.